Amendments to the Specification:

Please replace the paragraph beginning on page 2, at line 6 with the following rewritten paragraph:

Stereo vision techniques are commonly used in multiple camera systems to recover spatial information of the scene. Such systems yield a 3D range image where the range values may not be defined at every pixel. Imaging systems that are capable of recovering range values at every pixel (full 3D range recovery) are known in the art. For example, Cyberware, Inc. manufactures a system whereby a laser is scanned across a scene. Another method described in U.S. Patent ~~4,953,616~~ <u>4,935,616</u> (and further described in the Sandia Lab News, vol. 46, No. 19, September 16, 1994) provides a scannerless range imaging system using either an amplitude-modulated high-power laser diode or an array of amplitude-modulated light emitting diodes (LEDs) to completely illuminate a target scene. An improved scannerless range imaging system that is capable of yielding color intensity images in addition to the 3D range images is described in commonly-assigned, copending U.S. Patent Application Serial No. 09/572,522, <u>now U.S. Patent No. 6,349,174,</u> filed May 17, 2000 and entitled "Method and Apparatus for a Color Scannerless Range Imaging System". As used herein, a scannerless range imaging system will be referred to as a "SRI camera" and such a system is used in producing both intensity and 3D panoramas.

Please replace the paragraph beginning on page 3, at line 3 with the following rewritten paragraph:

Because of the nature of the SRI system, there is a further problem that must be addressed when merging two adjacent range images. The SRI system actually yields phase values that describe the phase offset for each pixel relative to one wavelength of the modulated illumination. These phase values are then converted to range values (because the modulation frequency is known). This leads to two types of ambiguity. First, if the objects in the scene differ in distances greater than one wavelength of the modulated illumination, the computed range values will reflect discontinuities where the corresponding phase values transitioned

from one cycle to the next. This ambiguity problem can be solved by the method described in commonly-assigned, copending U.S. Patent Application Serial No. 09/449,101, now U.S. Patent No. 6,288,776, which was filed November 24, 1999 in the names of N.D. Cahill et al. and entitled "Method for Unambiguous Range Detection). Even if the first type of ambiguity is resolved, a second type of ambiguity exists. This ambiguity arises because the phase values returned by the SRI system do not contain any information about absolute distance to the camera. The information captured by the SRI system is only sufficient to generate relative range values, not absolute range values. Therefore, the absolute range values differ by the values computed and returned by the SRI system in the range images by some unknown constant. In general, the unknown constant for a given range image is not the same as the unknown constant for another range image. This presents a problem when attempting to merge/stitch two adjacent range images captured from the SRI system. If the unknown constants are not the same, it will be impossible to continuously merge the two images.

Please replace the paragraph beginning on page 3, at line 26 with the following rewritten paragraph:

Therefore, two problems emerge. The first problem is that the computed 3D values in a given image are not absolutely known; they are only known relative to the other objects in the same image. Thus, an unknown constant offset must be added to every 3D value in the image. However, the constant offsets in subsequent 3D images may be different, and the difference in offsets must be determined in order to correctly merge the 3D values from neighboring scenes. Even if the first problem is solved, the 3D values of an object point in subsequent images are still dependent on orientation of the camera optical axis for each image. Consequently, distortion appears when a sequence of 3D images is used to describe the shape of an object. For instance, a smooth surface object in the three-dimensional space appears as a fragmented smooth surface object after reconstruction, using the untreated 3D images. Three methods have been shown to address the second problem in

panoramic 3D map formation. Each method comprises transforming 3D values into some reference coordinate system. As described in commonly assigned, copending U.S. Patent Application Serial No. 09/383,573, now U.S. Patent No. 6,507,665, filed August 25, 1999 in the names of Nathan D. Cahill and Shoupu Chen, and entitled "Method For Creating Environment Map Containing Information Extracted From Stereo Image Pairs", a directional transformation transforms 3D values by projecting points orthographically into a reference plane. As also described in Serial No. 09/383,573, a perspective transformation transforms 3D values by projecting points to the common nodal axis. As described in commonly assigned, copending U.S. Patent Application Serial No. 09/686,610, filed 11 October 2000 in the names of Lawrence A. Ray and Shoupu Chen, and entitled "Method for Three Dimensional Spatial Panorama Formation", an (X,Y,Z,) transformation transforms 3D values into 3-element vectors describing orthographic range to a reference system .

Please replace the paragraph beginning on page 11, at line 29 with the following rewritten paragraph:

The image processing method 100 forms a complete three-dimensional scene panorama for virtual reality visualization. The method 100 uses an image bundle 102 to generate a corresponding spatial image, e.g. an (X,Y,Z) image, in step 104. An inquiry of whether all image bundles have been captured is performed 106. A negative response to the inquiry causes the SRI camera to move to an adjacent position in step 108. A warping function and registration point is computed 110 and used to determine the differences in constant offsets of the relative 3D range values between image bundles captured from adjacent positions in step 112. Once these differences have been determined, they are applied to the spatial images in step 114. An arbitrary reference three-dimensional world coordinate system is established in step 116 to uniquely describe the spatial property of the scene captured. All the estimated spatial images are transformed in step 118 to the reference three-dimensional world coordinate system with a homogeneous transformation matrix

that is constructed based on the information of the capturing device. The transformed spatial images are stitched together to form a spatial panorama after a cylindrical warping procedure 120 and a registration process 122. Likewise, the intensity images are stitched together to form an intensity panorama in step 124 after the same procedures. Both spatial and intensity panoramas are used in a virtual display with no further transformation operation needed.

Please replace the paragraph beginning on page 12 at line 17 with the following rewritten paragraph:

The notion of an image bundle is an important aspect of a preferred range estimation method using an SRI camera. As shown in relation to Figure 2, an image bundle 200 includes a combination of images captured by the SRI system as well as information pertinent to the individual images and information common to all the images. The image bundle contains two types of images: range images 202 related to the image capture portion of the SRI process and an intensity image 204, which may be a color image. Common information 206 in the image bundle 200 would typically include the number of range images in the bundle (three or more) and the modulation frequency used by the SRI system. Other information might be the number of horizontal and vertical pixels in the images, and/or data related to camera status at the time of the image capture. Image specific information will include the phase offset 1...N used for each (1...N) of the individual range images 202. The image bundle 200 includes a minimum of three such images, each of which are monochrome. The additional intensity image 204 is an image using an optical channel of the SRI camera that does not contain range capture components. For example, as disclosed in the aforementioned ~~Serial No. 09/572,522~~ U.S. Patent No. 6,349,174, which is incorporated herein by reference, a beamsplitter is used to establish two optical paths: one path contains the range imaging elements and the other path contains regular optics for transmitting the intensity (e.g., color) image. An optical network (including light control means such as a shutter) recombines the image paths toward a

single image responsive element, and a range image and a intensity image are separately, and sequentially, captured. Alternatively, the range imaging elements and regular optics may be interchanged in a single optical path. Although the intensity image may be a color image, it is preferably, but not necessarily, the same size as the range images 202.

Please replace the paragraph beginning on page 13 at line 17 with the following rewritten paragraph:

One such warp that corrects for the distortion (but not the only such warp) is a cylindrical warp 110, where the images are warped onto a cylinder 304 about the vertical axis of the cylinder. This warping technique is described in detail in the aforementioned copending U.S. Patent Application Serial No. 09/383,573, now U.S. Patent No. 6,507,665, "Method For Creating Environment Map Containing Information Extracted From Stereo Image Pairs", which is incorporated herein by reference. Briefly described, the warp can be described by a function $W(x_p,y_p)$ that maps pixel 324 (x_p,y_p) in the image plane 318 to pixel 312 (x_c,y_c) in the warped plane 310. The cylindrical warping function $W(x_p,y_p)$ can be determined in the following manner; suppose the real world point 306 is projected through the rear nodal point 308 of the taking lens onto the cylinder 304 at point 312 (x_c,y_c), where x_c is the horizontal pixel coordinate 314 and y_c is the vertical pixel coordinate 316 (relative to the orthogonal projection of the nodal point 308 onto the image plane 318). The intensity/range value assigned to the cylindrically warped image at point 312 (x_c,y_c) should be the intensity/range value found at point 324 (x_p,y_p) in the planar image 318, where x_p is the horizontal pixel coordinate 320 and y_p is the vertical pixel coordinate 322 of point 324. It can be shown that (x_p,y_p) can be computed in the following way:

$$x_p = \frac{\tan(x_c p_x / f)}{p_x / f}, \qquad \text{(Eq. 10)}$$

$$y_p = \begin{cases} \dfrac{y_c \tan(x_c p_x / f)}{x_c p_x / f}, & x_c \neq 0 \\ y_c, & x_c = 0 \end{cases}, \qquad \text{(Eq. 11)}$$

where p_x is the length of pixels of the image plane 318 in the x-direction and f is the focal length of the taking lens. In general, (x_p,y_p) will not be integer valued, so

it is appropriate to interpolate nearby intensity values. For range values it is only appropriate to assign the value of the pixel nearest (x_p, y_p).

Please replace the paragraph beginning on page 15 at line 23 with the following rewritten paragraph:

Figure 6 describes the process 112 (referred to in Figure 1), whereby an estimate for α is determined. In 600, an initial estimate for α is chosen; e.g., $\alpha = 0$. In 602, the right hand side of Equation 13 is evaluated, yielding $\hat{d}_2$, an estimate of the 3D range value in the right image. In 604, the 3D range images are warped according to the warp function *W*, and then they are registered using the pre-determined registration point 400. The error between the predicted $\hat{d}_2$ values and the actual d_2 values in the overlap region 406 of the warped registered images are computed 606 by calculating the difference $\hat{d}_2 - d_2$ at each pixel, squaring this difference, and then summing the squared difference values for all overlapping pixels. An inquiry is made 608 as to whether the errors (measured by the summed squared difference values) is acceptable. If the result of the inquiry is negative, a new estimate for α is chosen according to some optimization scheme 610 (e.g., Newton's method, line search, etc., see Fletcher, *Practical Methods of Optimization*, 2nd Edition, John Wiley & Sons, 1987). A good choice is the Levenberg-Marquardt optimization scheme, which is described in the aforementioned Fletcher reference (pages 100-107). When the result of the inquiry 608 is finally affirmative, the current estimate for α is chosen 612 to be the relative range difference between the two images. According to 114 (referring to Figure 1), the relative range difference α is added to each 3D range value in the left image 500. Note that once the relative range difference has been applied, the range values in the left 500 and right 502 images will not be absolute; rather, they will still be relative, but with consistent constant offsets.